

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventorship..... Venkatesan et al.
Applicant..... Microsoft Corporation
Attorney's Docket No. MS1-650US
Title: Minimum Delta Generator for Program Binaries

TRANSMITTAL LETTER AND CERTIFICATE OF MAILING

To: Commissioner of Patents and Trademarks,
Washington, D.C. 20231

From: Kasey C. Christie (Tel. 509-324-9256; Fax 509-323-8979)
Lee & Hayes, PLLC
421 W. Riverside Avenue, Suite 500
Spokane, WA 99201

The following enumerated items accompany this transmittal letter and are being submitted for the matter identified in the above caption.

1. Specification—title page, plus 55 pages, including 22 claims and Abstract
2. Transmittal letter including Certificate of Express Mailing
3. 8 Sheets Formal Drawings (Figs. 1-9)
4. Return Post Card

Large Entity Status [x] Small Entity Status []

Date: 11/14/00

By: Daniel L. Hayes
Daniel L. Hayes
Reg. No. 34,618

CERTIFICATE OF MAILING

I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL685270435

Date: 11/14/00

By: Lori A. Vierra
Lori A. Vierra

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

Minimum Delta Generator for Program Binaries

Inventors:

Ramarathnam Venkatesan

Saurabh Sinha

ATTORNEY DOCKET NO. MS1-650US

TECHNICAL FIELD

This invention generally relates to generating a minimum delta between program binaries.

BACKGROUND

One of the important characteristics of modern software systems is its ability to be upgraded, which may be called "upgradability." Old software is continuously being replaced by newer versions, and code reusability and modular development are major features of software design.

Accuracy

When software is upgraded from an old version to a new version, complete accuracy is vital. Every bit in the newly upgraded software in the target computer must match exactly with the new software from at its media source. Otherwise, the new software may operate incorrectly or not at all.

To assure complete accuracy, conventional techniques completely replace the old software with the new software. As software programs (in particular major application suites and operating systems) grow in size and complexity, this wholesale replacement-to-update scheme becomes more time consuming and frustrating to the customer of such software.

Aggravating matters is a trend to move the source of such updates from local, portable, high-bandwidth removable media (such as a CD-ROM) to remote, centralized, relatively low-bandwidth network servers (such as Internet web servers). While replacing a 100 MB of software may be from a CD-ROM may

1 take several minutes, replacing the same amount of software over a dial-up
2 Internet connection may take several hours.

3 Herein, “complete accuracy” and “substantially identical” allows for minor
4 and insubstantial differences between the new software as originally produced and
5 the new software as it exists on a user’s computer.

6 7 **Conventional Delta-Patching**

8 Typically, newer versions of software have a few additional portions, as
9 well as some minor changes in older portions. Therefore, the brute force approach
10 of completely replacing the old with the new is overkill. An alternative is to
11 capture these changes into a “patch” so that one can reconstruct the newer version
12 from the older one. Because there are differences between the old and new
13 versions, this technique is sometimes called “diff-patching.” Herein, the
14 differences between the old and new versions are called the “delta” (Δ), thus this
15 diff-patching technique may be called “delta-patching” (or “ Δ -patching”).

16 The problem with delta-patching is accuracy. Identifying what is and is not
17 patched is difficult. If the boundaries of such a patch are not accurately
18 determined, the patched version will be different from the desired new version of
19 the software.

20 As a result, conventional delta-patching compromises efficiency to achieve
21 accuracy. Generally, the sub-module files, data files, library files, and groups of
22 such files are marked if there is any change whatsoever within them. This means,
23 for example, if one line of source code is changed within a 100 Kb DLL (dynamic
24 link library) file is changed, the entire DLL file is replaced. This is done rather
25 than replacing the fragment in the existing DLL file in part because of the

1 difficulty in selecting the fragment that needs replacing and replacing only that
2 with complete accuracy. However, it is most done because replacing the entire
3 module is more efficient with conventional techniques. A small change in one
4 little fragment might appear to be a change spread all over the entire program.

5 Although this conventional inefficient delta-patching is more efficient and
6 faster than wholesale replacement of the entire software, it still is not as efficiency
7 as possible. It would be more efficient to patch only those fragments of modules
8 and sub-modules that are different from or non-existent in the old software
9 version. Examples of fragments include subroutines, functions, objects, data
10 structures, interfaces, methods, or portions of any of these.

11 **Invariant Fragment Detection**

12
13 A prerequisite for detecting fragment deltas is the ability to detect
14 *invariance* of fragments. In other words, before a program module can be patched,
15 one needs to determine which fragments have not been changed across the two
16 versions. With knowledge of the source code for each version, detecting such
17 invariants and creating a patch is not very difficult.

18 However, detecting invariance of fragments becomes much more difficult
19 when dealing with binary manifestations of such fragments (with no knowledge of
20 the source code). A major difficulty is the existence of functionally unchanged
21 code that appears different in the differing versions of a program module. Code
22 may undergo no change in its functionality, but it may look different in the two
23 versions due to a variety of reasons. Examples of such reasons include:

- 24 • Changes in one region of code can cause another (unchanged) region
25 to look different

- Two small sequences of binary code may look identical even if they correspond to source code with different functionality
- Differences in the register allocation in the two builds

Change Begets Apparent Change

Often small changes in one portion of the code cause a cascade of changes in nearby and sometimes even far-off regions of code. Consider, for example, the following two source fragments:

Program P1	Program P2
-----	-----
function f(int p)	function g(int p)
int a=3, b=4;	int b=4, a=3;
if (b > p) {	if (b > p) {
a = p;	a = p;
return a;	return a;

The two functions f and g , located in the two programs $P1$ and $P2$, are really the same, apart from a difference of names. Clearly, knowledge of the source code would establish that the “if ($b > p$)” conditional in each fragment is the same, and need not be patched. However, if their corresponding binaries are examined, the offset of b from the base of the stack would be different in these two fragments. This is because of the declaration of a before b in $P1$ differs in form from the declaration of b before a in $P2$. Hence, the binaries of the two fragments will not be identical, even if everything else was the same. Of course, these differences in form are irrelevant in substance, but their resulting binaries are different nevertheless.

Now consider the following snippets:

```

Program P1
-----
x = f(10)

```

```

Program P2
-----
x = g(10)

```

Assume, for this example, that the functions f and g are defined as in the previous example. Here again, the two calls are identical, because the functions being called as well as the call arguments are identical. However, if the identity of f and g is not known, then the identity of the calls above will also not be discovered. This is an example of how local changes can cascade through potentially far-off regions of code.

Appear Identical, but Are Not

At times, two binary fragments may look identical even though they correspond to different regions in the structure of the corresponding programs. Consider the following:

<pre> Program P1 ----- int a = atoi(argv[1]); int b = atoi(argv[2]); if (a < 10) return; if (b < 20) return; </pre>	<pre> Program P2 ----- int b = atoi(argv[2]); if (b < 10) return; ... </pre>
--	---

The conditionals “if ($a < 10$)” in $P1$ and “if ($b < 10$)” in $P2$ might both translate to the same binary code, even though their functionality is different (as is seen clearly by examining their source code). This happens because the offset of b on the stack in $P2$ may be the same as that of a on the stack of $P1$. The two variables are clearly different, being defined by different program arguments, as

1 can be see in the context above them. However, comparing their binary
2 equivalents without reference to the source code context above can give the
3 illusion of an identity. A representation of the binary equivalent might look
4 something like:

```
5         mov eax, dword ptr[ebp+8h]  
6         cmp eax, 0ah  
7         jge L  
8         ret  
9         L: ...
```

9 **Register Allocation**

10 Another problem in detecting identity of binary fragment is caused by
11 *register allocation*. A change in a portion of code may cause the register allocation
12 to change in nearby regions, even though these latter regions have not been
13 modified. Therefore, when comparing binaries, one has to consider the possibility
14 that what looks like a change of register operands may in fact be an identity
15 disguised by a simple renaming of registers.

17 **SUMMARY**

18 Described herein is a technology for generating a minimum delta between
19 at least two program binaries. An implementation, described herein, is given a
20 source program (S) in a binary format and a target program (T) in a binary form.
21 It constructs control flow graphs (CFGs) of each. It matches common blocks of
22 the S's CFGs and T's CFGs. The blocks are matched based upon their content and
23 their local neighborhoods (e.g., *d*-neighborhoods). In addition, blocks are matched
24 using labels, which are based upon computed hash values. The matching is done
25 in multiple passes where each pass improves the matching by relaxing the criteria

1 for a match. In addition, the register renaming problems is solved so that blocks
2 can be fairly compared.

3 This described implementation produces an intermediate output, which is
4 the content of unmatched blocks. Such unmatched blocks are those found in T
5 that are not found in S. It generates a set of edge edit operations for merging the
6 unmatched blocks into S. The combination of the unmatched blocks and the edit
7 operations is the delta. To patch S to produce a reconstructed copy of T, the delta
8 is merged with S.

9 This summary itself is not intended to limit the scope of this patent. For a
10 better understanding of the present invention, please see the following detailed
11 description and appending claims, taken in conjunction with the accompanying
12 drawings. The scope of the present invention is pointed out in the appending
13 claims.

14 **BRIEF DESCRIPTION OF THE DRAWINGS**

15
16 The same numbers are used throughout the drawings to reference like
17 elements and features.

18 Fig. 1 is a schematic block diagram showing an embodiment of a minimum
19 delta generator for program binaries.

20 Fig. 2 is a schematic block diagram showing another embodiment of a
21 minimum delta generator for program binaries.

22 Figs. 3A-C are illustrations of a simplified application of an
23 implementation of a minimum delta generator for program binaries.

24 Figs. 4 and 5 are illustrations of applications of implementations of a
25 minimum delta generator for program binaries.

1 Fig. 6 is a flow diagram showing a methodological implementation of a
2 minimum delta generator for program binaries.

3 Fig. 7 is a flow diagram showing another methodological implementation
4 of a minimum delta generator for program binaries.

5 Fig. 8 is a flow diagram showing still another methodological
6 implementation of a minimum delta generator for program binaries.

7 Fig. 9 is an example of a computing operating environment capable of
8 implementing an implementation of a minimum delta generator for program
9 binaries.

10 **DETAILED DESCRIPTION**

11
12 The following description sets forth specific embodiments of a minimum
13 delta generator for program binaries that incorporate elements recited in the
14 appended claims. These embodiments are described with specificity in order to
15 meet statutory written description, enablement, and best-mode requirements.
16 However, the description itself is not intended to limit the scope of this patent.

17 Described herein are one or more exemplary implementations of a
18 minimum delta generator for program binaries. The inventors intend these
19 exemplary implementations to be examples. The inventors do not intend these
20 exemplary implementations to limit the scope of the claimed present invention.
21 Rather, the inventors have contemplated that the claimed present invention might
22 also be embodied and implemented in other ways, in conjunction with other
23 present or future technologies.

Incorporation by Reference

The following co-pending patent applications, which were all filed June 30, 1999 and assigned to the Microsoft Corporation, are incorporated by reference herein:

- U.S. Patent Application Serial No. 09/343,805 entitled "Translation and Transformation of Heterogeneous Programs";
- U.S. Patent Application Serial No. 09/343,298 entitled "Instrumentation and Optimization Tools for Heterogeneous Programs";
- U.S. Patent Application Serial No. 09/343,279 entitled "Shared Library Optimization for Heterogeneous Programs";
- U.S. Patent Application Serial No. 09/343,276 entitled "Application Program Interface for Transforming Heterogeneous Programs"; and
- U.S. Patent Application Serial No. 09/343,287 entitled "Cross Module Representation of Heterogeneous Programs".

Brief Overview

An exemplary implementation of a minimum delta generator for program binaries may be referred to as an exemplary "delta-generator." The one or more exemplary implementations described herein may be implemented (whole or in part) by a delta-generator system 100 of Fig. 1 and/or by a computing environment like that shown in Fig. 9.

To promote efficiency and speed, an exemplary delta-generator finds the minimum differences between the binary source and target programs and uses that

1 to modify the source program. However, accuracy is not compromised for such
2 efficiency and speed.

3 To accomplish this, the exemplary delta-generator compares the binary
4 fragments of the source and target programs. It identifies those fragments that are
5 the same (i.e., unchanged) between them. Necessarily, this also identifies those
6 fragments that are different (i.e., changed) between them. The changed fragments
7 are included in a delta. The exemplary delta-generator patches the source with the
8 delta to transform the source into the target program.

9 **High-Level Description of Delta-generator System**

10
11 Fig. 1 illustrates the delta-generator system 100. It shows a binary source
12 program 112 and a binary target program 122. The delta-generator system 100
13 reconstructs the target program 122 from the source program 112 and a delta (Δ)
14 142. Presumably, the target and source programs are different versions of the same
15 program, so that at least some portions of their code are common. For example,
16 the target program may be a newer version of the source program.

17 A source CFGer 110 produces a Control Flow Graph (CFG) of the source
18 program 112. CFGs are discussed in more detail below in the “Terminology”
19 section. Likewise, a target CFGer 120 produces a CFG of the target program 122.
20 Of course, these two CFGers 110 and 120 may be a common CFGer.

21 Fig. 1 shows a block matcher 130 that receives the source and target CFGs
22 as input and compares them. It matches binary fragments found in both the source
23 and target CFG.

24 The delta generator 140 identifies the remaining (i.e., unmatched)
25 fragments in the target CFG that were not matched as the delta (Δ) fragments.

1 Collectively, the block matcher 140 and the delta generator 150 it identify binary
2 fragments in the target CFG that are matched in the source CFG or those that are
3 unmatched in the source CFG.

4 In addition, the Δ -determiner 140 also determines how to edit the source
5 CFG to merge the Δ therein. These “edits” are part of the Δ 142, which is the
6 result of the Δ -determiner 140.

7 A Δ -patcher 150 patches the source program at the binary level. The source
8 program 112 is combined with the Δ 142 to reconstruct the target program 150.
9 The reconstructed target program 150 is identical to target program 122. It is
10 identical because the intended result is an identical copy of the target program 122.
11 Since this delta-generator system 100 introduces no inaccuracies by its efficient
12 patching of the source program, the copy will be identical.

13 Fig. 2 schematically illustrates a sample application scenario for the
14 exemplary implementation of a minimum delta generator for program binaries.
15 Fig. 2 shows a server side 210 and a client side 220 linked via a network
16 connection 215, such as the Internet. The server side 210 includes a server 212
17 having a copy of the binary source program (S) 112 and a copy of the binary target
18 program (T) 122. The client side 220 includes a client 222 having only a copy of
19 the binary source program (S) 112. Initially, it does not have a copy of the target
20 program (T) 122 or delta (Δ) 142.

21 The server 212 at the server side 210 produces the Δ 142, in accordance
22 with implementations described herein. It transmits such Δ to the client 222 at the
23 client side 220. Representations 142a and 142b (and their arrows) indicates the Δ
24 being transmitted from the server to the client.

1 The client patches S 112 with the newly arrived Δ 142 to reconstruct T 122.
2 Thus, with an implementation of an exemplary delta-generator, a minimum sized
3 delta is transmitted over a presumably limited bandwidth network 215 for a client
4 to patch S to accurately reconstruct T.

5 **Terminology**

7 An example of a program binary (or binary program) is binary
8 manifestation of program module (e.g., software). It does not include the source
9 code for a program module.

10 A “basic block” (or simply a “block”) is a sequence of instructions that has
11 a single entry point (only the first instruction can be reached from outside) and has
12 one exit point, at the last instruction. Thus, there is sequential control flow inside a
13 basic block. A basic block is an example of a fragment. Specifically, it is an
14 example of a code-fragment. A basic block may be called a “code block.”

15 A Control Flow Graph (CFG) is an abstraction of a program. It is a directed
16 graph where the nodes are the basic blocks in the program, and the edges (often
17 visually indicated by arrows) represent all possible control flow between blocks.
18 Figs. 3A-5 include CFGs.

19 Contiguous areas of static data in a program are called “data blocks” and
20 they also form nodes in the CFG. A data block is an example of a fragment.
21 Specifically, it is an example of a data-fragment.

22 Data blocks have static data (called the “raw” data) and they may have
23 pointers to other parts of the program. These pointers could be, for example,
24 virtual table entries of an object. Such pointers are relocatable data, and are thus
25 different from other static data in a data block.

1 In a CFG, pointers in data blocks are represented by directed edges from
2 the data block to the target blocks. In addition, address operands in the instructions
3 in a code block are represented as edges from the code block to the block at the
4 corresponding address.

5 The content of a data block includes its raw data, but it excludes the
6 pointers in the block. Likewise, the content of a code block includes its sequence
7 of instructions, but it excludes the address operands. The contents of the basic
8 blocks in a CFG, along with its edges, fully specify the program. Given any block
9 v in a CFG P , its parents are the set $\{x \mid x \rightarrow v \text{ is an edge in } P\}$, and its children are
10 the set $\{x \mid v \rightarrow x \text{ is an edge in } P\}$.

11 The CFG described herein contains complete information about the layout
12 of the program. Corresponding to the functions in the binary, there are procedures
13 in the CFG—each procedure is the subgraph induced by the basic blocks of a
14 single function. The layout of the function (i.e., the organization of the basic
15 blocks in the program's address space) is captured by a linked list of the all the
16 blocks of the corresponding procedure, such that a traversal of the list, from head
17 to tail, exactly describes the order of the basic blocks in the function's layout.

18 A tool V creates the CFG described herein for any given program binary.
19 Moreover, V may provide an interface (to the CFG) that allows a large class of
20 modifications to be made to the program. These modifications include, but are not
21 restricted to, addition and deletion of basic blocks, modification of their contents,
22 and changes in the edges.

23 Those who are of ordinary skill in the art understand how to implement and
24 use a tool V as described herein. In addition, the incorporated-by-reference patent
25 applications describe components that may be used to implement a tool like tool V .

Moreover, the following publications provide a general background on CFG methodologies that the tool *V* may use: Aho, Hopcroft, Ulman: "Principles of Compiler Design" and 2) A. Aho, R. Sethi, J. Ullman, "Compilers, Principles, Techniques, and Tools" (1986).

In describing the one or more examples of implementations of the minimum delta generator for program binaries, a source (*S*) binary of a program and a target (*T*) binary of a program are discussed. The relationship of *T* and *S* may be described mathematically as:

- $\text{Diff}(S, T) = \Delta$; and
- $\text{Patch}(S, \Delta) = T$

Or in other words as:

- The difference between *S* and *T* is delta; and
- When *S* is patched with delta, it is *T*.

As described herein, the exemplary implementations fully minimize Δ while maintaining the ability to achieve complete accuracy in patching *S* to reconstruct *T*. Herein, with regard to patching, "complete accuracy" (and "substantially identical" or similar terminology) allows for insubstantial differences between the original *T* and the reconstructed *T*.

Matching

A prerequisite for detecting changes between binary programs (such as *S* and *T*) is the ability to detect *identity* of binary fragments. In other words, before a program module can be patched, one needs to determine which fragments have not been changed across the two binary programs. For example, one may detect the identity of a fragment in a program module by finding the same unchanged

1 fragment in another program module (presumably an earlier or later version of the
2 same module).

3 4 Simple Example

5 Figs. 3A-3C illustrate a simple example of a methodological
6 implementation of the exemplary delta-generator. In this example, the initial goal
7 is to compute a Δ 330 for the input lists source (S) 310 and target (T) 320.

8 Fig. 3A shows source 310 and it includes six nodes, 311-316; and the
9 contents of those nodes are H, O, O, V, E, R, respectively. Target 320 includes six
10 nodes, 321-326; and contents of those nodes are H, O, O, P, E, R, respectively.
11 Compared to the source 310, the target 320 is largely unchanged. The difference
12 is node 314 (containing "V"), which is missing in the target but is replaced by
13 node 324 (containing "P").

14 The exemplary delta-generator matches nodes that have not changed. Based
15 on content and relative positioning, each node is matched (i.e., identified) except
16 for the nodes 314 and 324. For example, node 311 and 321 match because their
17 content (specifically, "H") and relative position are identical. Likewise, nodes 312
18 and 312 match, nodes 313 and 313, nodes 315 and 315, and nodes 316 and 316.

19 As shown in Fig. 3A, the delta 330 includes the content ("P") of node 324
20 as the content of node 337. The delta 330 also includes "edits" that specify how
21 the new node is patched into the source. Specifically, delta 330 specifies
22 "ADDEDGE(313, 337); ADDEDGE(337, 315)" within edit-box 339.

23 Fig. 3B shows that there are two edges to be added to source 310 after
24 inserting the new node 337. Thus, the delta contains the edit operations
25 ADDEDGE(313, 337), which translates to adding an edge 339a from the node 313

1 in the source 310 to the new node 337, and ADDEDGE(337, 315), which specifies
2 the addition of link 339b from node 337 to node 315 in the source.

3 Fig. 3C illustrates a combination of the source and the delta at 340a. The
4 addition of the delta (new node 337 and the two new edges (339a and 339b))
5 transforms the source to an exact copy of list target 320. That exact copy of the
6 target may be called a reconstructed target 340b.

7 Comparing Figs. 3B and 3C, one notices the deletion of edge 314a between
8 nodes 313 and 314 of the source and the deletion of edge 314b between nodes 314
9 and 315 of the source. This deletion is implicit by the addition of edges 339a and
10 339b.

11 Finding Matching Blocks

12
13 A good matching of blocks is one that minimizes the number of edit
14 operations. If two blocks are matched based on content similarity alone, there
15 might not be a good match. This is because there are usually multiple blocks in the
16 source as well as target with identical content, differing only in their location in
17 the respective graphs. For example, consider the subgraphs shown in Fig. 4.

18 Fig. 4 shows subgraphs of source 410 and target 420. In terms of content,
19 assume, for this example, that all the blocks in each of the following groups are
20 identical: (412, 416, 422), (413, 417, 423), (414, 418, 424), (411, 421), (415, 425).
21 Also, assume, for this example, that two blocks can be matched only if they are
22 identical in content. If block 422 is matched with block 412 (which both contain
23 "A"), then the best matching of the other blocks is

- 24 • block 411 (D_s) with block 421 (D_t);
- 25 • block 415 (E_s) with block 425(E_t);

- block 413 with block 423 (which both contain “B”); and
- block 414 with block 424 (which both contain “C”).

In this case, no edit operations are needed, since the local edge structure is preserved under this matching. On the other hand, if block 422 is matched with block 416 (which both contain “A”), the best matching of the other blocks is

- block 411 (D_s) with block 421 (D_t);
- block 415 (E_s) with block 425(E_t);
- block 417 with block 423 (which both contain “B”); and
- block 418 with block 424 (which both contain “C”).

The edit operations are Delete(411, 412), Add(422, 416), Delete(416, 419), Add(416, 415). Herein, Delete(a, b) means “delete the edge from node a to node b.” Similarly, Add(a, b) means “add an edge from node a to node b.”

The above examples are based upon content alone. These examples do not consider the blocks surrounding each block. Those surrounding blocks may be called the block’s neighborhood or local neighborhood. These examples illustrate that the conventional matching approaches that are based upon content alone do not produce a good matching of blocks. It may even produce incorrect matching.

Consequently, the exemplary delta-generator considers neighborhoods of blocks while matching them.

Neighborhood Consideration

The exemplary delta-generator matches blocks based on their *content and neighborhood*. In fact, it makes several passes of the source and target CFGs, considering neighborhoods of decreasing sizes in each pass. Matching based on larger neighborhood sizes first is desirable—it produces matches that are more

1 accurate. In a final pass, blocks are matched based on their content alone, after
2 matching based on neighborhood has been examined. Although not a requirement,
3 matching content last is advantageous because the overhead of recording an
4 unmatched block is usually greater than the size of the edit operations required
5 after it has been matched with any arbitrary block that has identical content.

6 To detect identical content of blocks, begin by obtaining is the sequence of
7 instruction opcodes for code blocks and the raw, non-relocatable data for data
8 blocks. As mentioned earlier, address operands are not considered part of content,
9 because addresses usually change even though the referred blocks remain same.
10 This problem is handled by considering address operands as edges in the CFG.

11 Register operands may be tricky to handle. Addition or deletion of some
12 blocks may change the register allocation in some other nearby regions of the
13 CFG, even though no change was made in these regions. Such operands are
14 matched modulo a possible register-renaming, using register flow analysis of each
15 graph. Since offsets of stack variables usually change because of changes in the
16 local variable declarations, they present a problem. To address this problem, the
17 exemplary delta-generator matches blocks even if certain immediate operands
18 have changed. In each of the above cases, when the sequence of operand types
19 remains same, the content of code blocks is considered unchanged.

20 An i -neighborhood (i.e., local neighborhood) of a node v in a CFG P is the
21 set of blocks that are at distance $\leq i$ from v in the undirected graph corresponding
22 to P . Thus, for example, a 0-neighborhood is the set $\{v\}$ and a 1-neighborhood is v
23 and its parents and children. The exemplary delta-generator computes, while
24 attempting to find a match for a block, its d -neighborhood, for different values of
25

1 d. Let a child edge for block v be an edge from v to one of its children, and a
2 parent edge be an edge from a parent of v to v .

3 The exemplary delta-generator computes the i -neighborhood of v doing a
4 breadth first traversal of the CFG, starting from the node v . It traverses child edges
5 as well as parent edges (in the reverse direction) of blocks until it is at a distance
6 greater than i from v . Alternatively, the exemplary delta-generator may ignore all
7 parent edges, and traverses only child edges to compute the neighborhood.

8 For instance, consider the first block v in a procedure f in S that is called
9 from three different blocks, b_1 , b_2 and b_3 . Suppose all these blocks, as well as the
10 procedure f remain unchanged in T , but an extra call is added to f from the block
11 b_4 . The i -neighborhood (for $i \geq 1$) of v changes because of this extra call, and thus,
12 v may remain unmatched because of it.

13 However, assume, for this example, that f itself did not undergo any
14 changes, and hence, it might be advantageous to ignore the parents of v while
15 computing its neighborhood. At times, the neighborhood may exclude any blocks
16 that are in a different procedure.

17 In general, the larger the neighborhood size considered, the more accurate
18 is the match, in terms of resolving ambiguities in matching candidates. However,
19 it leads to fewer matches, which is undesirable. Hence, The exemplary delta-
20 generator uses small neighborhood sizes (typically three or less), along with a
21 longer random component, which is obtained by doing a “random walk” that is
22 described later. This “random walk” heuristic has been empirically observed to
23 perform well.

24 One example of a “random walk” is illustrated by the following: Suppose a
25 person has to do a random walk on a straight road with several bus-stations

1 A,B....N (in that order) on it. The person starts at A (one end of the road). Since
2 this person can only go to the right, the person must go right – to B. Now, at B,
3 there are two options – go left (to A) or right (to C). In the uniform random walk
4 model, an unbiased coin is tossed and the person decides to go one way if its heads
5 and the other if its tails. If the person keeps doing this at every station she reaches,
6 then the person has performed a random walk.

7 The same idea can easily be extended to a general graph instead of a
8 straight line. Here is heuristic based upon the above illustration of “random walk”:
9 Let R be the undirected graph whose vertices are the nodes of the CFG and which
10 has an edge on a vertex pair if and only if the corresponding vertex pair in the
11 CFG has a (directed) control flow edge, in either direction. Clearly, R is a graph
12 with maximum degree three. It starts at the original node w , and at any node x ,
13 takes one of the dx edges with uniform probability. (dx is the degree of x). It
14 aborts when it encounters a procedure boundary (call or branch to another
15 procedure) or when the path length crosses a pre-determined limit.

16 Consider a block b that does not match any other block in content. Such a
17 block, if included in the neighborhood of any block v , may prevent v from being
18 matched, since in order for v to be matched to some block, its entire neighborhood
19 has to be matched. The block b is called a “bad” node or an outlier. Before
20 performing the matching, such outliers may be filtered out of the neighborhood of
21 all their neighbors.

22 The exemplary delta-generator detects matches in several passes of the
23 CFGs, based on progressively relaxing criteria. Preliminarily, the exemplary delta-
24 generator matches the procedures of S and T , based on a preliminary estimate of
25

1 the number of matching blocks. Examples of procedures are routines and sub-
2 routines of a binary program.

3 This preliminary matching is performed at a global level (i.e., blocks are
4 matched irrespective of the procedures in which they are located). Subsequently,
5 blocks are matched only if they belong to matching procedures. Such matching
6 may also be called local matching. The procedure match information is also used
7 to examine targets of call instructions to detect if two different calls target the
8 same function (which may have been renamed or whose first blocks might be
9 different).

10 During a matching pass, the current matching criteria are used to compute
11 short hash values of each node. The hash value of a block v is computed as
12 follows:

- 13 • Each block in its d -neighborhood (for some suitable value of d) is
14 hashed based on its content to produce a label for itself;
- 15 • All these labels are concatenated and the resulting string is hashed
16 again to produce a single label for v , which is called the “d-label” of
17 the block.

18 The nodes in each graph are sorted by these labels and blocks with identical
19 labels and matching procedures are matched.

20 Some other heuristics are used to improve the matching. These heuristics
21 are described herein, in particular in the “Additional Implementation Details”
22 sections below. Assume, for instance, two blocks c_1 and c_2 are said to be
23 *corresponding* children of blocks b_1 and b_2 if they are the i^{th} children of b_1 and b_2
24 respectively, according to the above ordering, for some i . An example of a
25 heuristic used to refine the matching is this: If matching blocks have

1 *corresponding children* that match in their content labels (but maybe not in their
2 *d*-labels), match those children.

3 At the end of this phase, the exemplary delta-generator has computed a
4 partial matching of the blocks of S and T.

5 Computing Edits

6
7 After the matching is complete, the next phase is computing the Δ . The toy
8 example of Figs. 3A-C gave a simple illustration of computing Δ . Fig. 5 illustrates
9 an example that shows how the matching computed above may be used in creating
10 the Δ for the case of graphs (i.e., CFGs).

11 Fig. 5 shows subgraphs of source 510 and target 520. Suppose the matching
12 phase has matched these block pairs: 511 and 521 (“A”), 512 and 522 (“B”), 513
13 and 523 (“C”), and 516 and 526 (“E”). The unmatched blocks in the target
14 subgraph 520 are block 525 (F_t) and block 524 (D_t). If Δ contains the contents of
15 block 525 (F_t) and block 524 (D_t), and the edit operations Add(512, 525) and Add
16 (512, 524), the target subgraph 520 can be reconstructed from the source subgraph
17 510.

18 The edit operations Delete(512, 514) and Delete(512, 515) need not be
19 explicitly included in the Δ . They are implicitly included. Those of ordinary
20 skilled in the art understand that the implicit edits may be expressly included in the
21 Δ without departing from the spirit and the scope of the present invention.
22 However, the implicit edits are not included in the Δ of the exemplary
23 implementations described herein because they unnecessarily increase the size of
24 the Δ .

1 The exemplary delta-generator, therefore, outputs (at some point) the
2 contents of the unmatched blocks in the target, and the edges that need to be
3 added, and that comprises the Δ .

4 **Methodological Implementation of the Exemplary Delta-generator**

5
6 Figs. 6 and 7 shows methodological implementations of the exemplary
7 delta-generator performed by the delta-generator system 100 (or some portion
8 thereof). These methodological implementations may be performed in software,
9 hardware, or a combination thereof.

10 Fig. 6 primarily shows the “matching” phase of a methodological
11 implementation of the exemplary delta-generator. Fig. 7 primarily shows the “edit”
12 phase of a methodological implementation of the exemplary delta-generator. The
13 “edit” phase follows the “matching” phase. The result of the “edit” phase is the
14 delta.

15 Fig. 6 shows, at 620, the exemplary delta-generator obtaining the sequence
16 of instruction opcodes for code blocks and the raw, non-relocatable data for data
17 blocks. Such blocks being blocks of procedures from source program binary (S)
18 612 and procedures from target program binary (T) 614.

19 At 624, the exemplary delta-generator computes, while attempting to find a
20 match for a block, its d -neighborhood, for different values of d . In addition, the d -
21 neighborhood is computed by breadth first traversal.

22 At 628, short hash values are computed for each node. In addition, the
23 exemplary delta-generator determines labels for nodes based upon the hash values
24 in their d -neighborhood (i.e., local neighborhood). Also at 628, the exemplary
25 delta-generator matches blocks even if certain immediate operands have changed.

At 630, the nodes in each graph are sorted by these labels, and blocks with identical labels and matching procedures are matched. At the end of this phase at 632, the exemplary delta-generator has computed a partial matching of the blocks of S and T. It also computes the converse, which are the unmatched blocks.

More details about the “matching” phase of this methodological implementation of the delta-generator are provided below in the “Additional Details” section.

Fig. 7 shows, at 720, the exemplary delta-generator receives the content 710 of unmatched blocks in P_t (i.e., a procedure from target program binary), which is part of the output from the methodological implementation of Fig. 6. The content of a block excludes all edge information. Thus, if an unmatched block v whose content is being recorded has address operands (for code blocks) or pointers (for data blocks), then such operands or pointers are replaced by “dummy” addresses that indicate that they have to be corrected using the edge edit information.

P_s and P_t represent all pairs of procedures (in S and T respectively) that are matched by the matching phase. For any pair, P_s is the procedure from the source and P_t is the matching procedure from T. The exemplary delta-generator receives a pair of procedure P_s 712 in source program binary (S) and procedure P_t 714 in source program binary (T).

At 730, all the information necessary to reconstruct a linked list of P_t from a linked list of P_s is recorded as part of the computed Δ . The layout (i.e., arrangement of blocks) of each procedure in the program’s address space is captured by a linked list of blocks, such that a sequential traversal of this list, from head to tail, can exactly describe the layout of the procedure. The first task in

1 reconstructing P_t from P_s would be to reconstruct the linked list of P_t from the
2 linked list of P_s . Hence, all the information necessary to reconstruct one linked list
3 from the other is recorded as part of the computed Δ .

4 At 732, edge edit operations are calculated. An edge edit operation, as
5 mentioned above, is an explicit addition of an edge (rather than an implicit
6 deletion). Each edge can be fully specified by its source and target. Within the
7 context of edits, the term “source” refers to the exact operand (for code blocks) or
8 relocatable word (for data blocks) that this edge represents. In addition, within the
9 context of edits, the word “target” on the other hand represents the block whose
10 address this operand or relocatable word refers to. The exemplary delta-generator
11 records an edited edge (an edge that has been added or deleted) by recording
12 unique identifiers for the source and the target in separate lists.

13 The delta for a given procedure (P_s) in the source is outputted at 734. Such
14 delta includes the unmatched blocks in P_t and the edge edit operations calculated
15 at 732.

16 The above methodological implementation of the “edit” phase of the
17 exemplary delta-generator records sufficient information to reconstruct the target
18 program binary up to some minor changes in register names, immediate values,
19 and instruction opcodes. These changed register names, immediate values, and
20 opcodes are recorded in separate lists.

21 Additional Implementation Details

22
23 The matching phase of the exemplary implementation makes several passes
24 of the CFGs, each of which is described in more detail below.
25

Multiple-Pass in the Matching Phase

Each pass in the matching phase uses different criteria for comparing blocks in S and T , and it does so by computing hash values or labels for each block. In the exemplary implementation, two exemplary sub-processes are employed and called in several passes. Herein, they are called `ComputeLabel` and `ComputeDLabel`.

ComputeLabel: Computes the hash value of the content of a block. This is also called the content label of the block. It uses a standard hash to calculate the hash value. An example of such a standard hash is MD5, which produces 16 byte strings as the hash of a given string of bytes.

This exemplary sub-process, called `ComputeLabel`, is parameterized to allow different matching criteria, and the Boolean parameters are: *Immediate*, *RegisterChain*, *OpcodeString* and *ProcId*. *Immediate* indicates whether the immediate operands in a code block are to be used in hashing. Similarly, *RegisterChain* specifies whether register renaming is to be used, and *OpcodeString* indicates whether an opcode is represented by its precise name or its group identifier (“group id”).

Using group ids for similar opcodes catches small changes in instructions that do not affect the functionality. For example, consider a portion of S where the only change is that an instruction “jge eax, 10” has changed to “jle eax, 10”. This could happen, for instance, if a bug was discovered where the direction of the inequality was incorrect. In such cases, both ‘jge’ and ‘jle’ can be represented by their opcode group ‘branch’ and the corresponding blocks might be matched as a result.

The fourth parameter, *ProcId* indicates whether procedure matching has already been done, in which case procedures have unique identifiers that match for matching procedures. If this parameter is set, function call targets are represented by their procedure identifiers. An example of the pseudo-code is for the sub-process ComputeLabel is:

```

ComputeLabel (Block, Immediate, RegisterChain, OpcodeString, ProcId)
If (Block is a data block)
    Iterate through each Relocatable word in Block, and clear it
    to zero.
    Hash the resulting data buffer and return the hash value.
Endif
If (Block is a code block)
    Create a buffer that may hold the data to be hashed.
    For each instruction except the last :
        If (OpcodeString is true)
            Record the opcode string in the buffer
        Else
            Record the opcode group in the buffer
        End
        For each operand in the instruction
            Case OperandType
                Pointer:    Record the constant string "Block".
                Register:   If (RegisterChain is true)
                            Record the register flow id and
                            register flow id is defined (computed earlier).
                            Else
                                Record the register name.
                            Endif
                Immediate:  If (Immediate is true)
                            Record the immediate constant
                            Endif
            End
        End
    End
End
For the last instruction :
    If (OpcodeString is true)
        Record the opcode string in the buffer
    Else
        Record the opcode group in the buffer
    End
    If (ProcId is true)
        Compute the target block of the instruction.
        If the target is in a different procedure
            Record the procedure id of the target procedure
        Endif
    Endif
End
Hash the buffer and return the hash value.
Endif

```

End ComputeLabel

ComputeDLabel: Computes the hash value based on content as well as neighborhood. This examines the entire already-computed neighborhood (including the root node) and concatenates their content labels. It then returns the hash value of this concatenated list.

The implementation uses a standard hash to calculate the hash value. An example of such a standard hash is MD5, which produces 16 byte strings as the hash of a given string of bytes.

Additional Details about the Methodological Implementation of the Matching Phase of the Exemplary Delta-generator

Fig. 8 illustrates the broad methodological implementation of the matching phase of the exemplary delta-generator. This broad methodological implementation may be performed by the delta-generator system 100 (or some portion thereof). This broad methodological implementation may be performed in software, hardware, or a combination thereof.

Additional details of this broad methodological implementation are described below:

Task 1 (810 of Fig. 8): Preliminary Matching of Basic Blocks

This task (810 of Fig. 8) matches basic blocks at a global level.

- For each block in S as well as T , the exemplary delta-generator calls ComputeLabel and assigns it a content label. Each of the four parameters to ComputeLabel is false at this stage.

- The exemplary delta-generator matches blocks of S with blocks of T based on the computed labels. Duplicate blocks in S (i.e., blocks having the same hash value) are matched to any of the blocks in T that have the same label. Whenever two blocks are matched, each of them is assigned a unique MatchId for reference later. Blocks that are unmatched after this task are designated as outliers.
- The exemplary delta-generator computes d -neighborhoods ($d=2$) for each block, by doing a breadth-first traversal that starts at this block. The traversal visits parents of a block if it has no children.
- The exemplary delta-generator filters out bad nodes computed above from all d -neighborhoods.
- The exemplary delta-generator calls ComputeDLabel for each block and assigns it a d -label, which is the value returned from the call.
- The exemplary delta-generator matches blocks of S and T that have identical d -labels.
- Any two matching blocks, b_1 and b_2 must have the same number of out-edges. Let these out-edges be ordered by their location in the blocks. Two blocks c_1 and c_2 are said to be corresponding children of blocks b_1 and b_2 if they are the i^{th} children of b_1 and b_2 respectively, according to the above ordering, for some i . The exemplary delta-generator refines the matches by using the following rule: If matching blocks have corresponding children that match in their content labels (but maybe not in their d -labels), match those children.

1 *Task 2 (812 of Fig. 8): Matching Procedures*

2 This task (812 of Fig. 8) uses the preliminary match computed above to
3 match procedures.

- 4 • The exemplary delta-generator iterates through each procedure P_t in
5 T , and computes, for each procedure P_s in S , the number m_{st} of
6 blocks in P_t that match blocks in P_s .
- 7 • The exemplary delta-generator matches procedure pairs based on
8 their m_{st} values. The m_{st} values may also be called the procedure-
9 matching-criteria. The technique for matching procedures using m_{st}
10 values is discussed below.
- 11 • The exemplary delta-generator assigns the same unique identifiers to
12 matching procedures. The exemplary delta-generator assigns unique
13 identifiers to all other procedures also.

14
15 An exemplary delta-generator for matching procedures meets this
16 condition: Given any P_s and any two P_{t1} and P_{t2} , if P_s is matched to P_{t1} then either
17 $m_{st1} \geq m_{st2}$ or P_{t2} is matched to P_s such that $m_{st2} \geq m_{st1}$. The following technique
18 satisfies the condition:

- 19 • Create (P_t, P_s, m_{st}) triplets $\forall s, t$ and sort them on m_{st} ; and
- 20 • Iterate through the sorted list of triplets in the descending order. At
21 any iteration, if P_t and P_s are both unmatched, match them.

22
23 *Task 3 (814 of Fig. 8): Local Matching of Basic Blocks*

1 This task (814 of Fig. 8) uses the procedure matching information obtained
2 in the previous task to compute local matching of blocks (i.e., blocks are matched
3 only between matching procedures).

- 4 • Start with no blocks matched. For each block in S as well as T , the
5 exemplary delta-generator calls `ComputeLabel` and assigns it a
6 content label. The *Immediate*, *OpcodeString* and *Procid* parameters
7 are true when calling `ComputeLabel`.
- 8 • The exemplary delta-generator matches blocks of S with blocks of T
9 based on the computed content labels. Duplicate blocks in S (i.e.,
10 blocks having the same hash value, are matched to any of the blocks
11 in T that have the same label). Whenever two blocks are matched,
12 each of them is assigned a unique `MatchId` for reference later. If a
13 block already has a previously assigned `MatchId`, reassign it. Blocks
14 that are unmatched after this task are designated as outliers.
- 15 • For $d = 3, 2, 1$ the exemplary delta-generator does the following:
 - 16 ○ Computes d -neighborhood for each block, by doing a
17 breadth-first traversal that starts at this block. The d -
18 neighborhood has a random component also. The random
19 component is a set of blocks that are encountered while
20 performing a fixed length uniform random walk starting at the
21 original node. The random walk looks only at real control
22 flow edges and chooses one of them with uniform probability.
23 It aborts when it encounters a procedure boundary (e.g., call
24 or branch to another procedure).

- Filters out the bad nodes computed above from all the d -neighborhoods.
- Calls ComputeDLabel for each block and assign it a d -label.
- Matches blocks of S and T that have identical d -labels. If two blocks in T match the same block in S , only one of them is actually matched. The tie may be broken arbitrarily.
- Improves the matching by following the guideline: If matching blocks have corresponding children that match in their content labels (but maybe not in their d -labels), match those children.

At any stage, a block previously matched is not rematched or unmatched.

- The exemplary delta-generator assigns each matching pair of blocks a unique MatchId, which is used as the content label for those two blocks thereafter.

Task 4 (816 of Fig. 8): Solving the Register Renaming Problem

This task (816 of Fig. 8) solves the register-renaming problem (i.e., it detects matching blocks even in the face of a possible register renaming). The exemplary delta-generator computes all register flow chains in S and T , and assigns each of them an identifier. Matching blocks have their corresponding register flows assigned the same identifier.

The exemplary delta-generator calls ComputeLabel with RegisterChain true and assigns each block a content label. The exemplary delta-generator matches unmatched blocks, based on these labels and their d -neighborhoods ($d=2$). The

1 exemplary delta-generator repeats this task until no more matches are discovered.
2 This is done because the new matching blocks may cause some more register
3 chains to be assigned the same identifiers, and in turn lead to even more matches.

4 Below, the section titled "Solving the Register Renaming Problem"
5 discusses this in more detail.

6 7 *Task 5 (818 of Fig. 8): The Final Pass*

8 In this last matching task (818 of Fig. 8), any pair of blocks with identical
9 content label (computed in the previous tasks) is matched.

10 11 Solving the Register Renaming Problem

12 For each procedure, the exemplary delta-generator computes its use-def
13 chains, def-use chains, and dominator information. A use-def chain is a list that
14 starts with a register use and goes through all the definitions (defs) of that register
15 that reach this use. Similarly, a def-use chain is a list that starts with a register def
16 and goes through each use of the same register, reached by this def. A block d is
17 called a dominator of another block b . b is reachable from d and any path to b
18 from the entry point of the procedure must first pass through d . Typically,
19 computing dominator information means computing all dominators of every block
20 in the procedure.

21 For each basic block in each procedure, the exemplary delta-generator
22 looks at all the register defs in the block, and assigns each of them an id, called the
23 def id, which is unique for that block. The exemplary delta-generator looks at all
24 the register uses in that block, and if any use has its only def in the same block,
25 assigns the def identifier to this use. Thus at this stage, all defs and some uses have

ids associated with them. There may be some register uses with undefined ids. ComputeLabel (defined above) when called with RegisterChain=true uses their names if the id is not defined.

The exemplary delta-generator calls ComputeLabel (with RegisterChain = true) for each unmatched block in S and T , and assigns the returned label to its content label. The exemplary delta-generator computes d -neighborhoods ($d=2$) for each unmatched block. The exemplary delta-generator calls ComputeDLabel to assign it a d -label, and matches S and T based on the newly computed d -labels.

For any pair of code blocks (b_s , b_t) currently matched, the exemplary delta-generator does the following:

- Iterates through corresponding (located at the same place) register defs in the two blocks.
- Assigns a unique identifier (positive), unique for the entire procedure, to each of the two defs being examined currently.

These identifiers replace any previously assigned def identifiers that the defs may have. The fact that two defs match is recorded, and a single identifier is associated with both of them.

Next, the exemplary delta-generator matches register uses based on the newly assigned def identifiers. For both S and T , the exemplary delta-generator iterates through all the register uses in each basic block. For each such use, all the defs are examined that reach it and are not dominated by it. The exemplary delta-generator adds up their def identifiers and assign the sum to the current use's identifier. Consequently, each register operand, def or use, has an identifier associated with it. Matching blocks have matching identifiers associated with corresponding register operands. More specifically, two register operands, that

1 have exactly matching defs, may also have matching identifiers even though the
2 register names may be different.

3 If any new matches were discovered in the most recent execution of the
4 exemplary delta-generator calling ComputeLabel with RegisterChain = true
5 (described above), the process loop back to there.

6 7 Additional Details about the Methodological Implementation of the Edit Phase of 8 the Exemplary Delta-generator

9 The exemplary delta-generator creates a new dummy CFG, called bin, with
10 a single procedure that may serve as a container for all the additional blocks of T .
11 These blocks are, for example, maintained in the form of a linked list.

12 The exemplary delta-generator traverses each block in T , in sequence, and
13 assigns consecutive identifiers (starting at zero) to each unmatched block. Also,
14 the exemplary delta-generator dumps the unmatched blocks into bin, making sure
15 that address operands in code blocks and pointers in data blocks are modified into
16 'dummy' addresses referring to the blocks containing those pointers.

17 The exemplary delta-generator traverses each block in S , in sequence, and
18 assigns consecutive identifiers to each matching block. The same identifiers are
19 assigned to matching blocks in T . The identifiers assigned in this and the previous
20 task are called New_Node_Ids.

21 The exemplary delta-generator examines the matched blocks of the source
22 graph. If a block x in S matches block y in T , check each of its outgoing edges.
23 For each edge $x \rightarrow z_1$, there has to be a corresponding $y \rightarrow z_2$.

- 24
25 • case i. z_2 matches z_1 : do nothing.

- case ii. z_2 matches $z_3 \neq z_2$: Record the link $x \rightarrow z_3$.
- case iii. z_2 is unmatched: Record the link $x \rightarrow z_2$.

An edge is recorded as follows. Three lists, named TargetId, OperandIndex and SourceProcFirstBlock are created. If an edge $x \rightarrow y$ is to be recorded, the New_Node_Ids of y is inserted into TargetId. If this edge was the i^{th} edge in the current procedure, i is pushed into the list OperandIndex. Finally, if any procedure P_s in S has at least one edge being recorded, the New_Node_Id of its first block is pushed into SourceProcFirstBlock.

The exemplary delta-generator examines all the unmatched blocks in the target graph. The exemplary delta-generator looks at each of its outgoing edges $x \rightarrow y$. If y is also an unmatched block, the exemplary delta-generator creates an edge in bin going from x to y . If y matches z in the source graph, the exemplary delta-generator records the link $x \rightarrow z$ by pushing the New_Node_Id of z into a list called TargetOperandTargets.

For each procedure P_s in S that has a matching procedure P_t in T , the exemplary delta-generator traverses the linked list of the blocks of P_t . At any point in the traversal, if ptr_1 and ptr_2 point to the current node and its next node in T , let mptr_1 and mptr_2 be their ‘matching’ pointers. If ptr_1 points to block x in T , its matching pointer mptr_1 points to its matching block in S if there is one, or to the copy of x that has been dumped in bin. At any stage in the traversal, if mptr_1 and mptr_2 point to nodes that are not adjacent (in that order) in the list, the edge $\text{mptr}_1 \rightarrow \text{mptr}_2$ is recorded. This is done by pushing the New_Node_Ids of their corresponding blocks into separate lists called LinkedListsLeft and LinkedListsRight. In addition, the New_Node_Id of the first block of P_t , along

1 with the latter's size (in terms of number of blocks) and the number of recorded
2 links in P_i is recorded in a separate list called ProcedureInfo.

3 The exemplary delta-generator uses the edit information computed above,
4 along with S to reconstruct T .

5 Alternatively, the exemplary delta-generator may detect the register names,
6 immediate values, and instruction opcodes that differ between the original and
7 reconstructed versions of T , and output these differences in a separate list called
8 MinorDifferences.

9 All these lists, together with the bin comprises the computed Δ in an
10 exemplary implementation of the delta-generator.

11 **Delta Compression**

12
13 The above sections sufficiently describe how to reconstruct, using the
14 exemplary delta-generator, a copy of T , given S and Δ . Like what is shown in Fig.
15 2, the generated Δ may be transmitted from a server 212 to a client 222.

16 This Δ may be compressed to use the available bandwidth on network 215
17 more efficiently. Therefore, the exemplary delta-generator may format the patch
18 in a way that it can be heavily compressed. The patch is split into several parts,
19 each holding information of a different kind, so that each part can be compressed
20 separately. This approach is based on the observation that different types of data
21 (code vs. lists of integers, for example) may be best compressed by different
22 compression algorithms (i.e., engines).

23 This approach may be combined with steps mentioned above in the
24 "Additional Implementation Details" section. In the exemplary delta-generator, the
25 "several parts" mentioned here are the several different lists that were mentioned

1 in the Edits phase above and the dummy “bin” that was mentioned there. Each of
2 these lists (and bin) are stored as a separate file in one implementation, and thus
3 compressed separately. An example of a compression algorithm that may be used
4 is LZW.

5 It is expected to work better for large executables (of the order of several
6 hundred kilobytes) and our initial experiments show that for such large
7 executables, The exemplary delta-generator produces, on the average, a smaller
8 patch that any of the other available patching tools that we know of.

9 10 **Examples of Applications for Exemplary Implementations**

11 Among the possible examples of applications for implementations of the
12 exemplary delta-generator is as a tool to send software upgrades or patches via
13 media with bandwidth constraints. With software availability on the Internet
14 becoming increasingly popular every day, such a patching tool may be beneficial
15 to uses of software.

16 Another possible example of applications for implementations of the
17 exemplary delta-generator is in the realm of anti-piracy of software. Using an
18 implementation of the exemplary delta-generator, an anti-piracy tool may be able
19 to detect similarity of programs. Another possible exemplary application could be
20 in analysis of the nature of changes between different versions of software.

21 22 **Exemplary Computing Environment**

23 Fig. 9 illustrates an example of a suitable computing environment 920 on
24 which an exemplary delta-generator are, for example, implemented.
25

1 Exemplary computing environment 920 is only one example of a suitable
2 computing environment and is not intended to suggest any limitation as to the
3 scope of use or functionality of an exemplary delta-generator. Neither should the
4 computing environment 920 be interpreted as having any dependency or
5 requirement relating to any one or combination of components illustrated in the
6 exemplary computing environment 920.

7 The exemplary delta-generator is operational with numerous other general
8 purpose or special purpose computing system environments or configurations.
9 Examples of well known computing systems, environments, and/or configurations
10 that are, for example, suitable for use with an exemplary delta-generator include,
11 but are not limited to, personal computers, server computers, thin clients, thick
12 clients, hand-held or laptop devices, multiprocessor systems, microprocessor-
13 based systems, set top boxes, programmable consumer electronics, network PCs,
14 minicomputers, mainframe computers, distributed computing environments that
15 include any of the above systems or devices, and the like.

16 An exemplary delta-generator are, for example, described in the general
17 context of computer-executable instructions, such as program modules, being
18 executed by a computer. Generally, program modules include routines, programs,
19 objects, components, data structures, etc. that perform particular tasks or
20 implement particular abstract data types. An exemplary delta-generator may also
21 be practiced in distributed computing environments where tasks are performed by
22 remote processing devices that are linked through a communications network. In
23 a distributed computing environment, program modules are, for example, located
24 in both local and remote computer storage media including memory storage
25 devices.

1 As shown in Fig. 9, the computing environment 920 includes a general-
2 purpose computing device in the form of a computer 930. The components of
3 computer 920 may include, by are not limited to, one or more processors or
4 processing units 932, a system memory 934, and a bus 936 that couples various
5 system components including the system memory 934 to the processor 932.

6 Bus 936 represents one or more of any of several types of bus structures,
7 including a memory bus or memory controller, a peripheral bus, an accelerated
8 graphics port, and a processor or local bus using any of a variety of bus
9 architectures. By way of example, and not limitation, such architectures include
10 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
11 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
12 local bus, and Peripheral Component Interconnects (PCI) bus also known as
13 Mezzanine bus.

14 Computer 930 typically includes a variety of computer readable media.
15 Such media are, for example, any available media that is accessible by computer
16 930, and it includes both volatile and non-volatile media, removable and non-
17 removable media.

18 In Fig. 9, the system memory includes computer readable media in the form
19 of volatile memory, such as random access memory (RAM) 940, and/or non-
20 volatile memory, such as read only memory (ROM) 938. A basic input/output
21 system (BIOS) 942, containing the basic routines that help to transfer information
22 between elements within computer 930, such as during start-up, is stored in ROM
23 938. RAM 940 typically contains data and/or program modules that are
24 immediately accessible to and/or presently be operated on by processor 932.
25

1 Computer 930 may further include other removable/non-removable,
2 volatile/non-volatile computer storage media. By way of example only, Fig. 9
3 illustrates a hard disk drive 944 for reading from and writing to a non-removable,
4 non-volatile magnetic media (not shown and typically called a "hard drive"), a
5 magnetic disk drive 946 for reading from and writing to a removable, non-volatile
6 magnetic disk 948 (e.g., a "floppy disk"), and an optical disk drive 950 for reading
7 from or writing to a removable, non-volatile optical disk 952 such as a CD-ROM,
8 DVD-ROM or other optical media. The hard disk drive 944, magnetic disk drive
9 946, and optical disk drive 950 are each connected to bus 936 by one or more
10 interfaces 954.

11 The drives and their associated computer-readable media provide
12 nonvolatile storage of computer readable instructions, data structures, program
13 modules, and other data for computer 930. Although the exemplary environment
14 described herein employs a hard disk, a removable magnetic disk 948 and a
15 removable optical disk 952, it should be appreciated by those skilled in the art that
16 other types of computer readable media which can store data that is accessible by a
17 computer, such as magnetic cassettes, flash memory cards, digital video disks,
18 random access memories (RAMs), read only memories (ROM), and the like, may
19 also be used in the exemplary operating environment.

20 A number of program modules are, for example, stored on the hard disk,
21 magnetic disk 948, optical disk 952, ROM 938, or RAM 940, including, by way of
22 example, and not limitation, an operating system 958, one or more application
23 programs 960, other program modules 962, and program data 964.

24 Each of such operating system 958, one or more application programs 960,
25 other program modules 962, and program data 964 (or some combination thereof)

1 may include an embodiment of an exemplary delta-generator. More specifically,
2 each may include an embodiment of delta-generator system, comparator, edit-op
3 determiner, and output sub-system.

4 A user may enter commands and information into computer 930 through
5 input devices such as keyboard 966 and pointing device 968 (such as a "mouse").
6 Other input devices (not shown) may include a microphone, joystick, game pad,
7 satellite dish, serial port, scanner, or the like. These and other input devices are
8 connected to the processing unit 932 through a user input interface 970 that is
9 coupled to bus 936, but are, for example, connected by other interface and bus
10 structures, such as a parallel port, game port, or a universal serial bus (USB).

11 A monitor 972 or other type of display device is also connected to bus 936
12 via an interface, such as a video adapter 974. In addition to the monitor, personal
13 computers typically include other peripheral output devices (not shown), such as
14 speakers and printers, which may be connected through output peripheral interface
15 975.

16 Computer 930 may operate in a networked environment using logical
17 connections to one or more remote computers, such as a remote computer 982.
18 Remote computer 982 may include many or all of the elements and features
19 described herein relative to computer 930.

20 Logical connections shown in Fig. 9 are a local area network (LAN) 977
21 and a general wide area network (WAN) 979. Such networking environments are
22 commonplace in offices, enterprise-wide computer networks, intranets, and the
23 Internet.

24 When used in a LAN networking environment, the computer 930 is
25 connected to LAN 977 via network interface or adapter 986. When used in a

1 WAN networking environment, the computer typically includes a modem 978 or
2 other means for establishing communications over the WAN 979. The modem
3 978, which are, for example, internal or external, are, for example, connected to
4 the system bus 936 via the user input interface 970 or other appropriate
5 mechanism.

6 Depicted in Fig. 9, is a specific implementation of a WAN via the Internet.
7 Computer 930 typically includes a modem 978 or other means for establishing
8 communications over the Internet 980. Modem 978, which are, for example,
9 internal or external, is connected to bus 936 via interface 970.

10 In a networked environment, program modules depicted relative to the
11 personal computer 930, or portions thereof, are, for example, stored in a remote
12 memory storage device. By way of example, and not limitation, Fig. 9 illustrates
13 remote application programs 989 as residing on a memory device of remote
14 computer 982. It will be appreciated that the network connections shown and
15 described are exemplary and other means of establishing a communications link
16 between the computers are, for example, used.

17 **Exemplary Operating Environment**

18
19 Fig. 9 illustrates an example of a suitable operating environment 920 in
20 which an exemplary delta-generator may be implemented. Specifically, the
21 exemplary delta-generator(s) described herein is implemented (wholly or in part)
22 by any program module 960-962 and/or operating system 958 in Fig. 9 or a
23 portion thereof.

24 The operating environment is only an example of a suitable operating
25 environment and is not intended to suggest any limitation as to the scope or use of

1 functionality of the exemplary delta-generator(s) described herein. Other well
2 known computing systems, environments, and/or configurations that are suitable
3 for use with an exemplary delta-generator include, but are not limited to, personal
4 computers (PCs), server computers, hand-held or laptop devices, multiprocessor
5 systems, microprocessor-based systems, programmable consumer electronics,
6 wireless phones and equipments, general- and special-purpose appliances,
7 application-specific integrated circuits (ASICs), network PCs, minicomputers,
8 mainframe computers, distributed computing environments that include any of the
9 above systems or devices, and the like.

10 **Computer-Executable Instructions**

11
12 An implementation of an exemplary delta-generator may be described in
13 the general context of computer-executable instructions, such as program modules,
14 executed by one or more computers or other devices. Generally, program modules
15 include routines, programs, objects, components, data structures, etc. that perform
16 particular tasks or implement particular abstract data types. Typically, the
17 functionality of the program modules may be combined or distributed as desired in
18 various embodiments.

19 **Computer Readable Media**

20
21 An implementation of an exemplary delta-generator may be stored on or
22 transmitted across some form of computer readable media. Computer readable
23 media can be any available media that can be accessed by a computer. By way of
24 example, and not limitation, computer readable media may comprise "computer
25 storage media" and "communications media."

1 “Computer storage media” include volatile and non-volatile, removable and
2 non-removable media implemented in any method or technology for storage of
3 information such as computer readable instructions, data structures, program
4 modules, or other data. Computer storage media includes, but is not limited to,
5 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
6 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
7 tape, magnetic disk storage or other magnetic storage devices, or any other
8 medium which can be used to store the desired information and which can be
9 accessed by a computer.

10 “Communication media” typically embodies computer readable
11 instructions, data structures, program modules, or other data in a modulated data
12 signal, such as carrier wave or other transport mechanism. Communication media
13 also includes any information delivery media.

14 The term “modulated data signal” means a signal that has one or more of its
15 characteristics set or changed in such a manner as to encode information in the
16 signal. By way of example, and not limitation, communication media includes
17 wired media such as a wired network or direct-wired connection, and wireless
18 media such as acoustic, RF, infrared, and other wireless media. Combinations of
19 any of the above are also included within the scope of computer readable media.
20
21
22
23
24
25

Conclusion

Although the minimum delta generator for program binaries has been described in language specific to structural features and/or methodological steps, it is to be understood that the minimum delta generator for program binaries defined in the appended claims is not necessarily limited to the specific features or steps described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed present invention.

1 **CLAIMS:**

2
3 **1.** A method for generating a delta between a first program binary and
4 a second program binary, the method comprising the steps of:

5 obtaining a first control flow graph (CFG) representation of the first binary
6 and obtaining a second CFG representation of the second binary;

7 comparing the first and second CFG representations to identify blocks
8 (nominally matched blocks) that match in the first and second CFG
9 representations, thereby identifying blocks (nominally unmatched blocks) in the
10 second CFG representation that do not match in the first CFG representation;

11 determining edit-operations that merges the unmatched blocks into the first
12 CFG representation so that first CFG representation is substantially identical to the
13 second CFG representation;

14 producing a delta comprising the unmatched blocks and the edit-operations.

15
16 **2.** A method for transmitting a delta comprising:

17 a method for generating a delta as recited in claim 1;

18 transmitting the delta over a network.

19
20 **3.** A method for patching a copy of the first program binary, the
21 method comprising:

22 a method for generating a delta as recited in claim 1;

23 patching the copy of the first program binary so that the copy is
24 substantially identical to the second program binary, wherein the delta guides such
25 patching.

1
2 4. A method as recited in claim 1, wherein the comparing step further
3 comprises matching blocks across the first and second CFG representations based,
4 at least partially, upon content of blocks being compared and neighborhoods of
5 blocks local to blocks being compared.

6
7 5. A method as recited in claim 1, wherein the comparing step further
8 comprises matching blocks across the first and second CFG representations based,
9 at least partially, upon content of the blocks being compared and a local
10 neighborhoods of blocks surrounding the blocks being compared, wherein a local
11 neighborhood of a block is augmented with a random sampling of blocks from a
12 substantially large neighborhood of blocks surrounding the block.

13
14 6. A computer-readable medium having embodied thereon a data
15 structure, comprising a delta generated in accordance with the steps recited in
16 claim 1.

17
18 7. A computer-readable medium having computer-executable
19 instructions that, when executed by a computer, performs the method as recited in
20 claim 1.
21
22
23
24
25

1 **8.** A method for matching blocks between a first control flow graph
2 (CFG) representation of a portion of a first program and a second CFG
3 representation of a portion of a second program, the method comprising:

4 matching blocks between the first and second CFG representations based
5 upon the content of the blocks;

6 detecting outliers, wherein outliers are blocks in the first CFG
7 representation that do not match any block in the second CFG representation
8 during the matching step;

9 computing a neighborhood of each block in the first and second CFG
10 representations by performing a breadth first traversal;

11 removing the outliers from each neighborhood.

12
13 **9.** A method as recited in claim 8 further comprising:

14 computing labels for each block in first and second CFG representations
15 based upon content of a block;

16 for each neighborhood computed in the computing step, forming a “d-
17 label” for each block in a neighborhood based upon the labels of the blocks within
18 the neighborhood;

19 d-label matching blocks between first and second CFG representations
20 based upon the d-label of the blocks.

21
22 **10.** A computer-readable medium having computer-executable
23 instructions that, when executed by a computer, performs the method as recited in
24 claim 8.
25

1 11. A method for matching procedures between a first control flow
2 graph (CFG) representation of a portion of a first program and a second CFG
3 representation of a portion of a second program, the method comprising:

4 computing a procedure-match-criterion for each procedure in the second
5 CFG representation, where the procedure-match-criterion for a procedure in the
6 second CFG representation represents the number of matching blocks between that
7 procedure and a specified procedure in the first CFG representation;

8 matching procedures in the second CFG representation with the specified
9 procedure in the first CFG representation based upon the procedure-match-criteria
10 for the procedures in the second CFG representation.

11
12 12. A computer-readable medium having computer-executable
13 instructions that, when executed by a computer, performs the method as recited in
14 claim 11.

15
16 13. A method for matching of blocks in a procedure of a first control
17 flow graph (CFG) representation of a portion of a first program between an
18 ostensibly matching procedure of a second CFG representation of a portion of a
19 second program, the method comprising:

20 matching blocks between the first and second CFG representations based
21 upon the content of the blocks;

22 computing successively smaller neighborhoods of each block in the first
23 and second CFG representations via breadth first traversals;

1 for each neighborhood computed in the computing step, forming a “d-
2 label” for each block in a neighborhood based upon the labels of the blocks within
3 the neighborhood;

4 d-label matching blocks between first and second CFG representations
5 based upon the d-label of the blocks.

6
7 **14.** A method as recited in claim 13, wherein at least one neighborhood
8 computed in the computing steps is augmented with a random sampling of blocks
9 in the complete representation of the neighborhood.

10
11 **15.** A computer-readable medium having computer-executable
12 instructions that, when executed by a computer, performs the method as recited in
13 claim 13.

14
15 **16.** A method for facilitating matching of blocks between a first control
16 flow graph (CFG) representation of a portion of a first program and a second CFG
17 representation of a portion of a second program, whereby blocks may match
18 except for nominal differences, the method comprising:

19 computing register flows for each block by analyzing a def-use chain and a
20 use-def chain;

21 for each block, assigning def-identifiers to definitions (“def’s”) of each
22 register;

23 for each block, associating a use of a register in a block with the def
24 identifier of all defs of that register that reach this use;
25

1 matching blocks between the first and second CFG representations based
2 upon the def-identifiers and use-identifiers.

3
4 **17.** A computer-readable medium having computer-executable
5 instructions that, when executed by a computer, performs the method as recited in
6 claim 16.

7
8 **18.** A patch data structure generated in accordance with the following
9 acts:

10 providing a server computer in a communications with a communications
11 network;

12 receiving input from a client computer by way of the communications
13 network, the input providing a parameter indicative of a request for upgrading a
14 copy of a first program binary to a match a second program binary;

15 retrieving a delta between a first program binary and the second program
16 binary, wherein computing such delta comprises the steps of:

17 a) obtaining a first control flow graph (CFG) representation of the first
18 binary and obtaining a second CFG representation of the second
19 binary;

20 b) comparing the first and second CFG representations to identify
21 blocks (nominally matched blocks) that match in the first and second
22 CFG representations, thereby identifying blocks (nominally
23 unmatched blocks) in the second CFG representation that do not
24 match in the first CFG representation;

1 c) determining edit-operations that merges the unmatched blocks into
2 the first CFG representation so that first CFG representation is
3 substantially identical to the second CFG representation;
4 d) producing a delta comprising the unmatched blocks and the edit-
5 operations;
6 generating the patch data structure as a function of the delta.

7
8 **19.** A method for transmitting a patch data structure comprising
9 transmitting a patch data structure as recited in claim 18 over a communications
10 network.

11
12 **20.** A method for patching a copy of the first program binary at a client
13 computer, the method comprising patching the copy of the first program binary so
14 that the copy is substantially identical to the second program binary, wherein a
15 delta in a patch data structure as recited in claim 18 guides such patching.
16
17
18
19
20
21
22
23
24
25

1 **21.** A delta-generator system, comprising:

2 a comparator that is configured to compare a first control flow graph (CFG)
3 representation of a first program binary and a second CFG representation of the
4 second program binary for identifying blocks (nominally matched blocks) that
5 match in the first and second CFG representations, thereby identifying blocks
6 (nominally unmatched blocks) in the second CFG representation that do not match
7 in the first CFG representation;

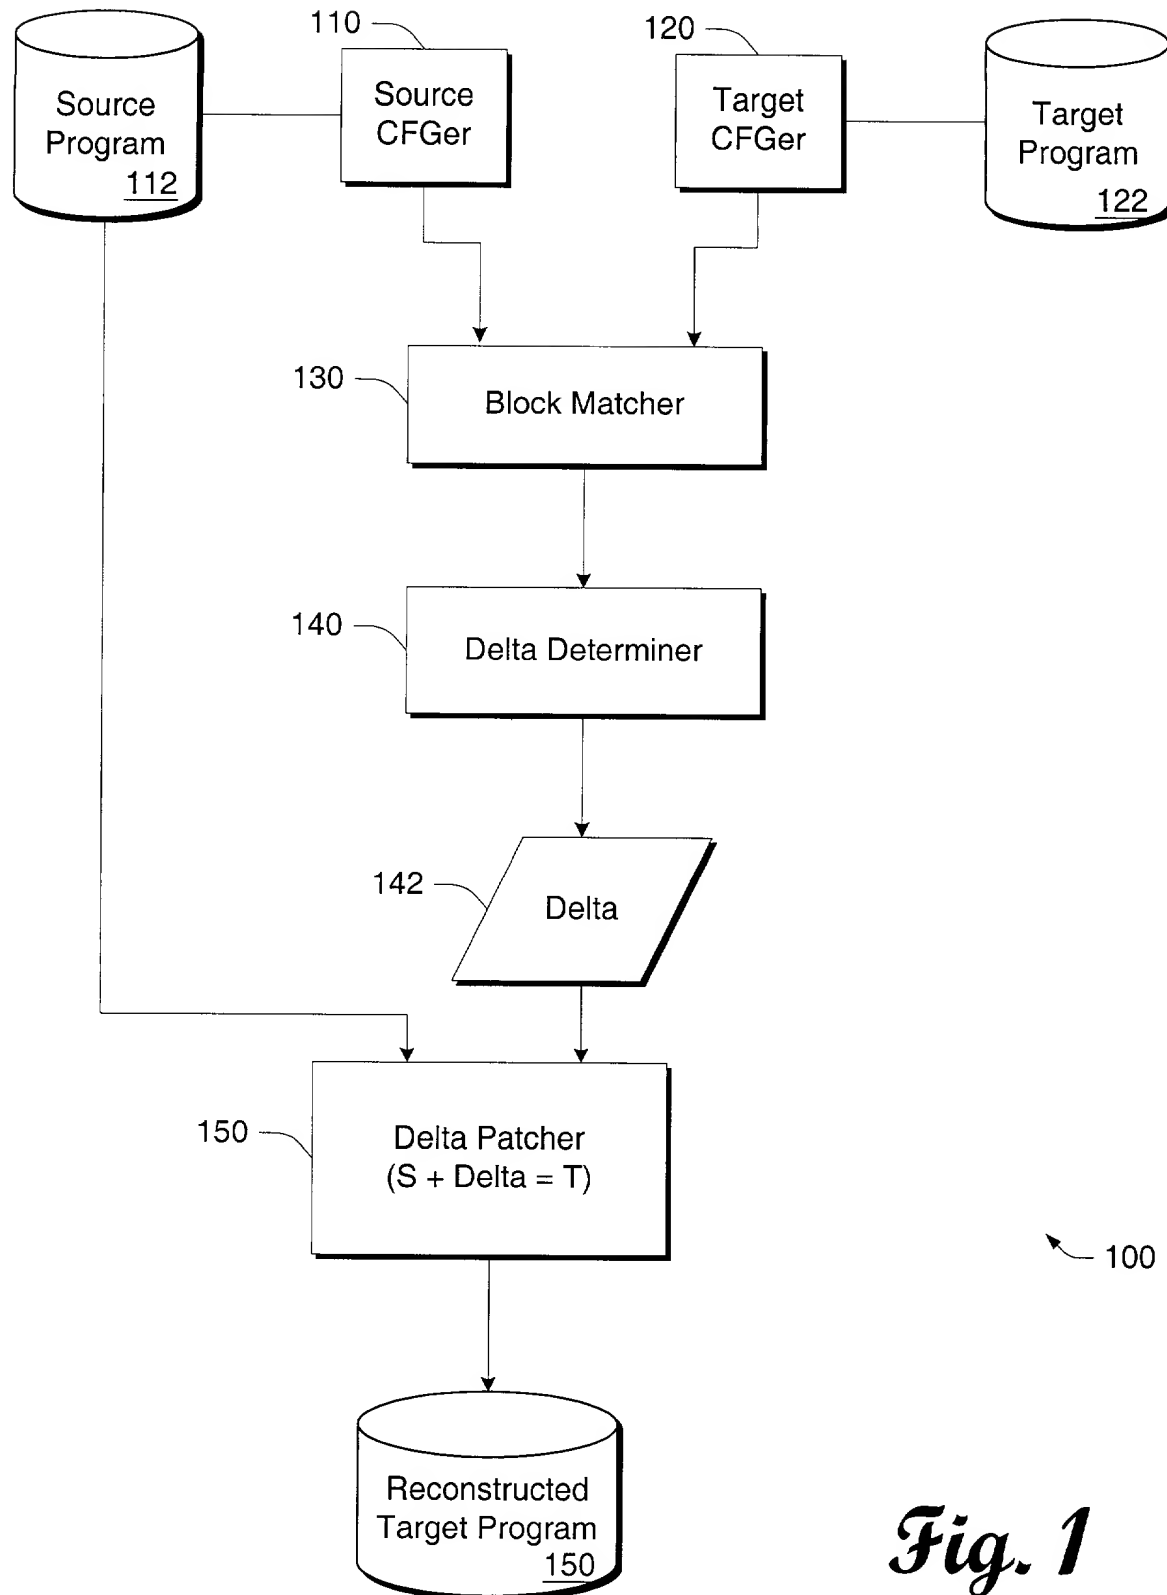
8 an edit-op determiner configured to determine the edit-operations that
9 merges the unmatched blocks into the first CFG representation so that first CFG
10 representation is substantially identical to the second CFG representation;

11 an output sub-system that is configured to produce a delta comprising the
12 unmatched blocks and the edit-operations.

13
14 **22.** A computer-readable medium having embodied thereon a data
15 structure comprising a delta produced by the system as recited in claim 21.

1 **ABSTRACT**

2
3 Implementations of this invention provide a technology for generating a
4 minimum delta between at least two program binaries. An implementation of this
5 invention is given a source program (S) in a binary format and a target program
6 (T) in a binary form. It constructs control flow graphs (CFGs) of each. It matches
7 common blocks of the S's CFGs and T's CFGs. The blocks are matched based
8 upon their content and their local neighborhoods. In addition, the register
9 renaming problems is solved so that blocks can be fairly compared. This
10 implementation of this invention produces an intermediate output, which is the
11 content of unmatched blocks. It generates a set of edge edit operations for merging
12 the unmatched blocks into S. The combination of the unmatched blocks and the
13 edit operations is the delta. To patch S to produce a reconstructed copy of T, the
14 delta is merged with S. This abstract itself is not intended to limit the scope of this
15 patent. The scope of the present invention is pointed out in the appending claims.
16
17
18
19
20
21
22
23
24
25

*Fig. 1*




Fig. 2

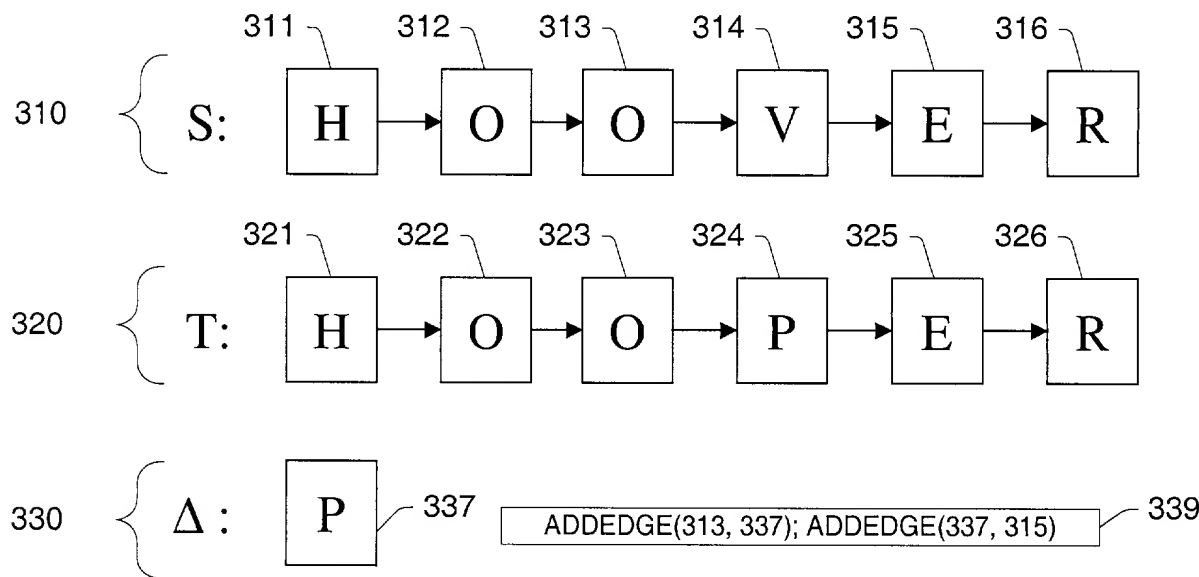

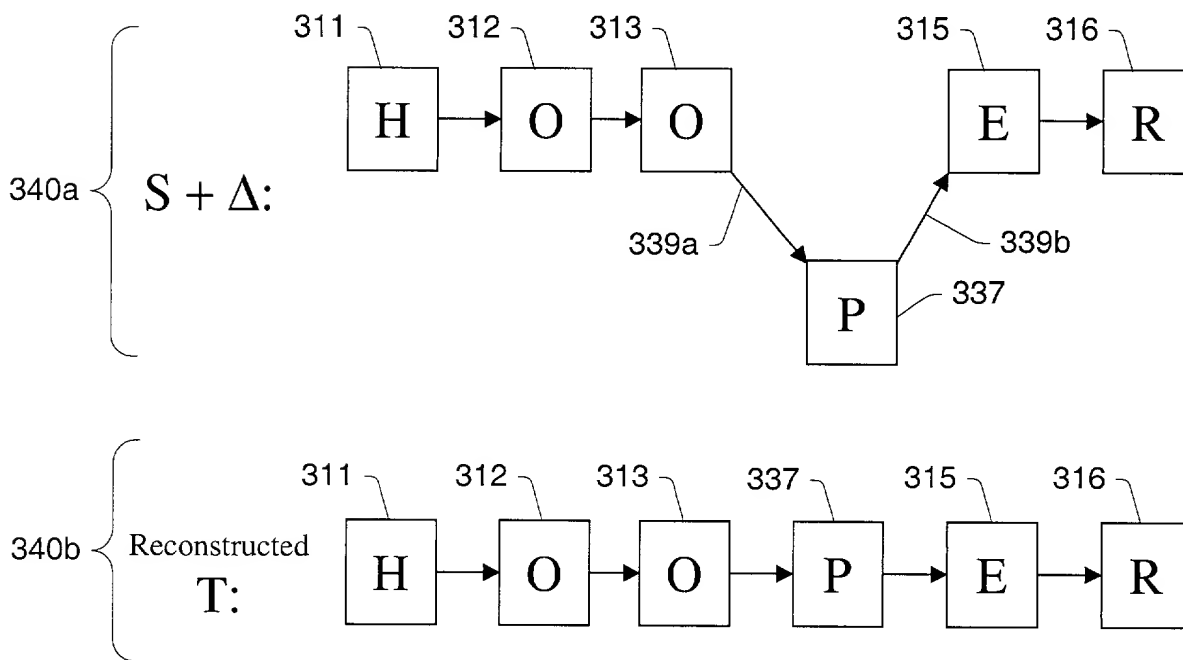
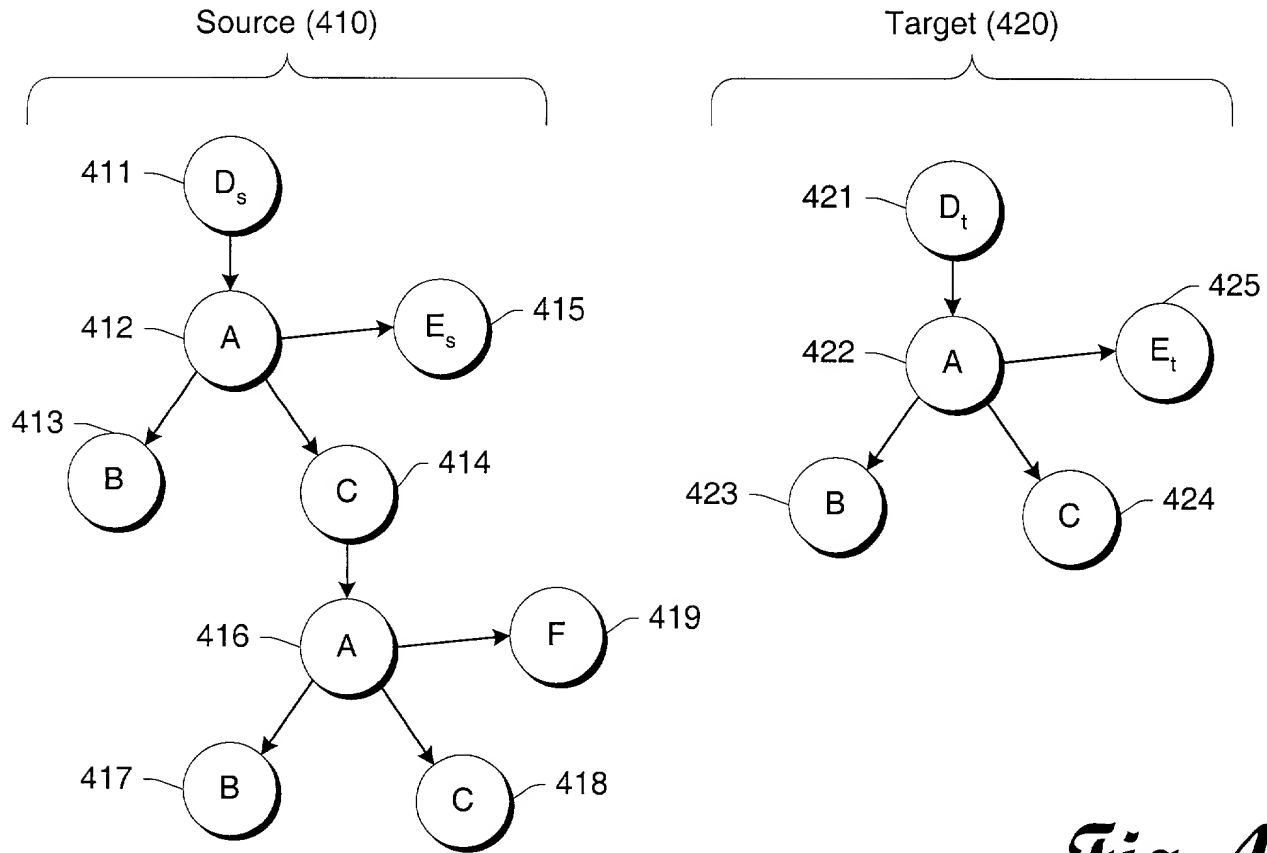

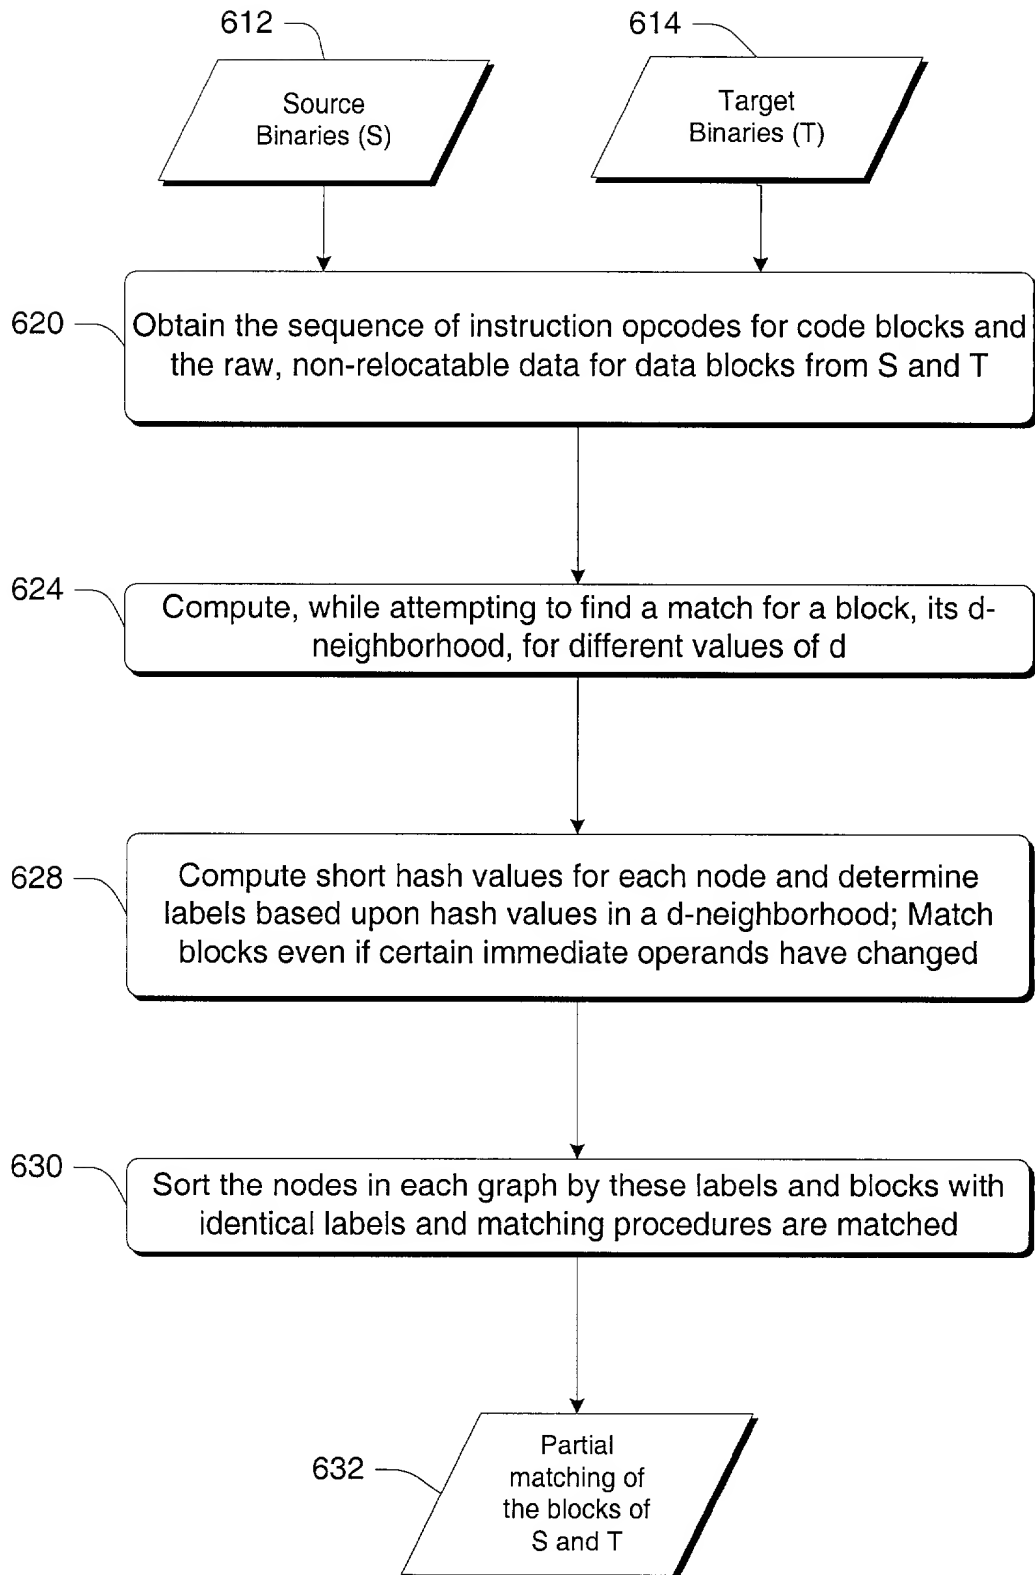
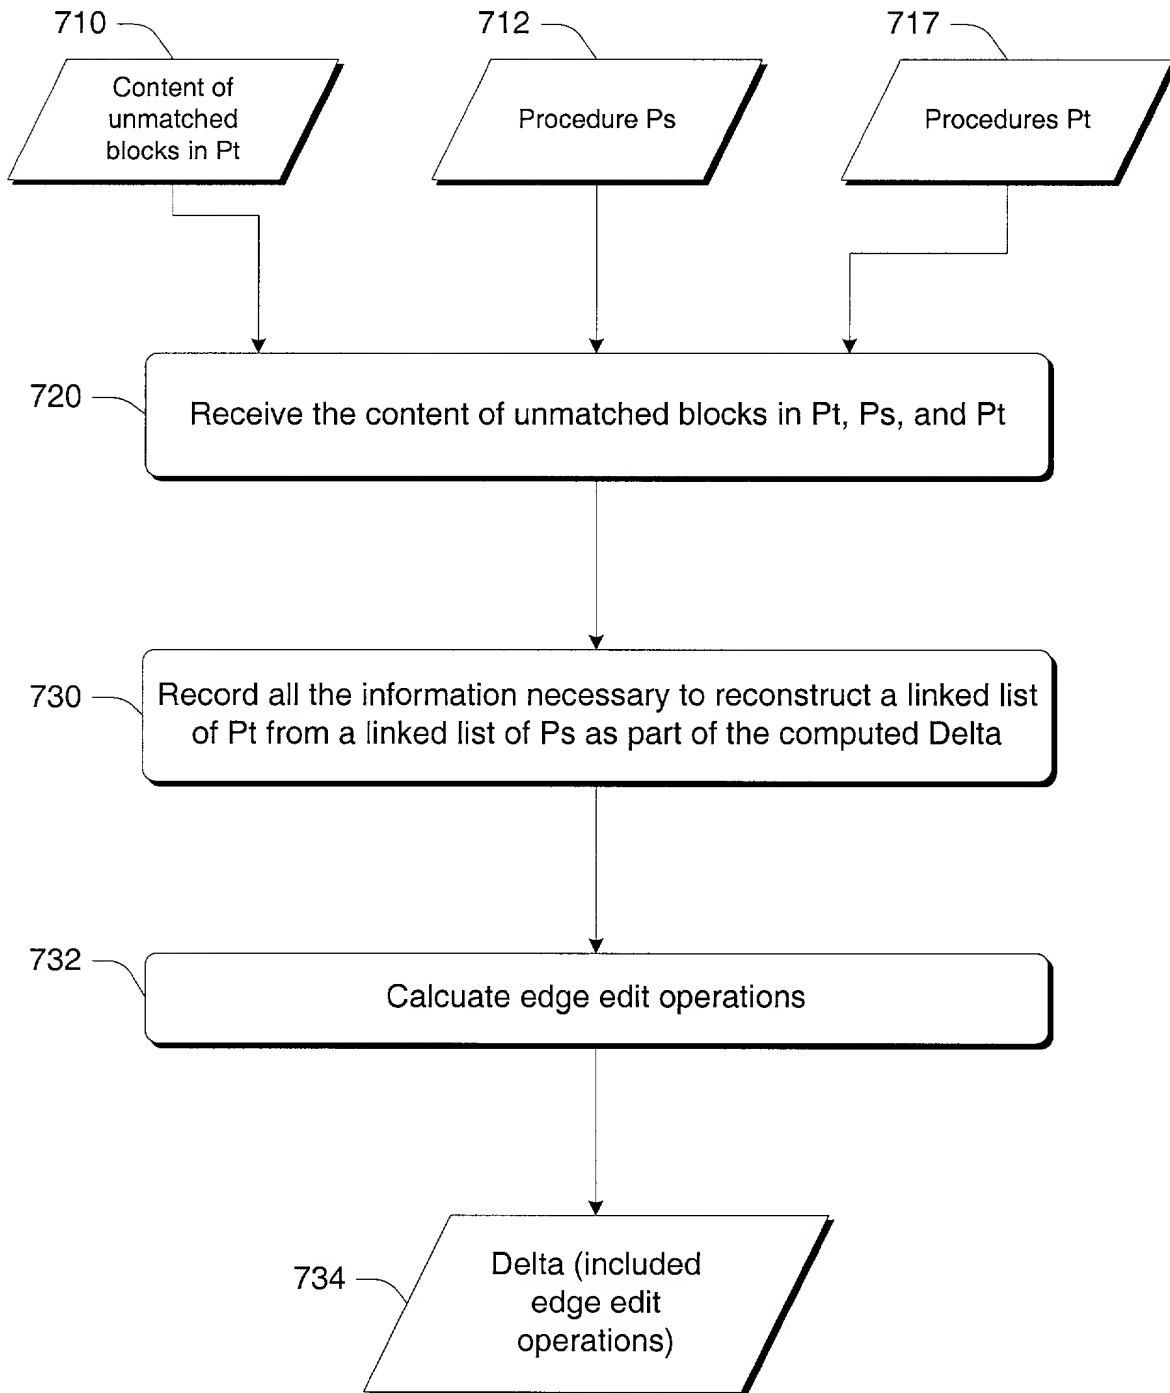


Fig. 3A

*Fig. 3B**Fig. 3C*

*Fig. 4**Fig. 5*

*Fig. 6*

*Fig. 7*

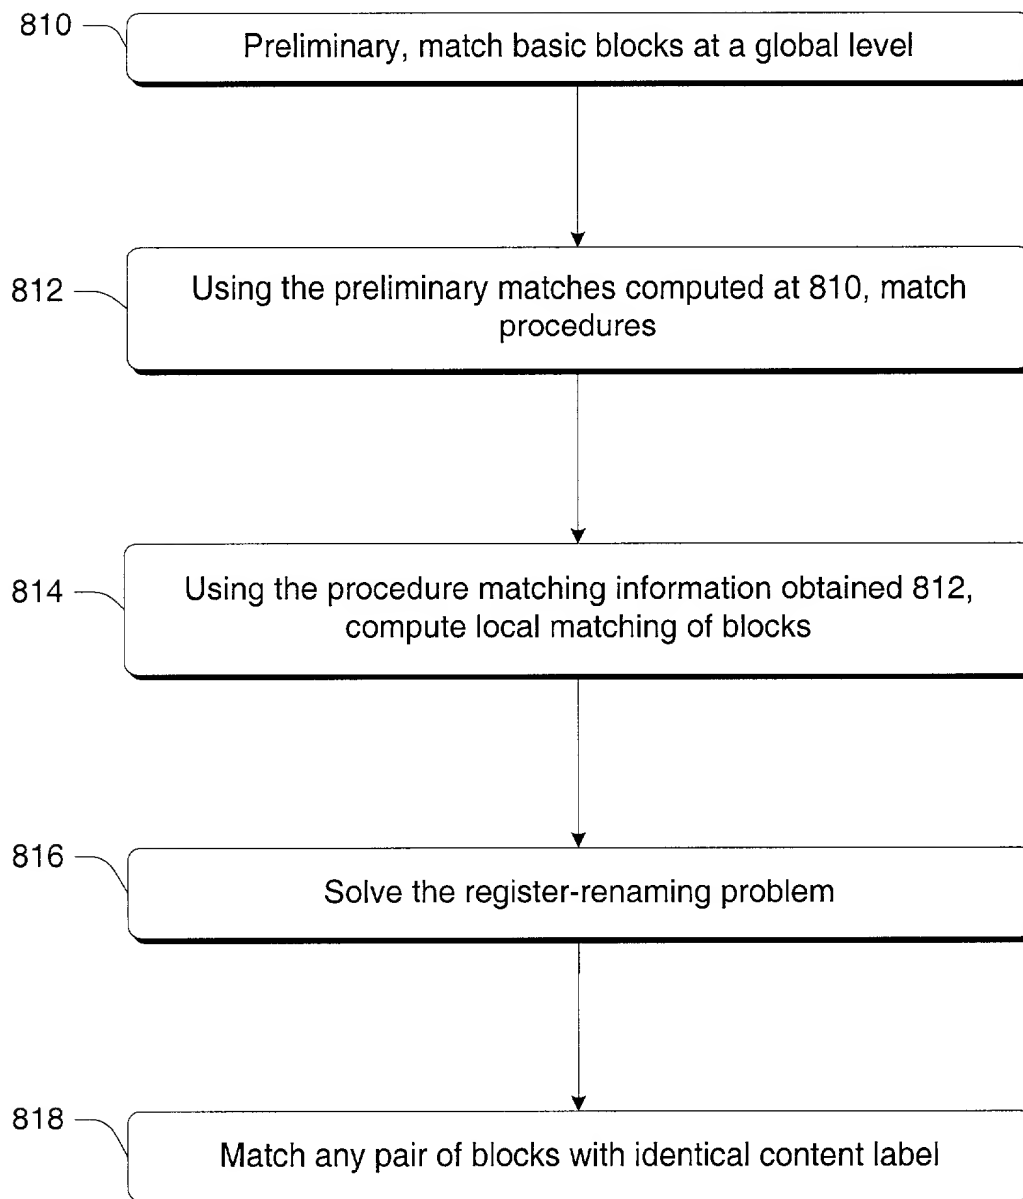
*Fig. 8*

Fig. 9

